Docket No.: 1509-489 **PATENT**

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of :

Christopher MARTIN :

New U.S. Patent Application :

Filed: February 23, 2004 :

For: DATA STORAGE DRIVE AND METHOD EMPLOYING DATA COMPRESSION

CLAIM OF PRIORITY AND TRANSMITTAL OF CERTIFIED PRIORITY DOCUMENT

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In accordance with the provisions of 35 U.S.C. 119, Applicant hereby claims, in the present application, the priority of British Patent Application No. 0304052.4, filed February 22, 2003. The certified copy is submitted herewith.

Respectfully submitted,

LOWE HAUPTMAN GILMAN & BERNER, LLP

Allan M. Lowe

Allan M. Lowe
Registration No. 19,641

1700 Diagonal Road, Suite 310
Alexandria, Virginia 22314
(703) 684-1111 AML/gmj
Facsimile: (703) 518-5499
February 23, 2004

The Patent Office

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

I also certify that the attached copy of the request for grant of a Patent (Form 1/77) bears an amendment, effected by this office, following a request by the applicant and agreed to by the Comptroller-General.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed HBehen

Dated 13 February 2004

Patents Form 1/77

Pat Act 1977
(Ru .5)

THE PATENT OFFICE
A
22 FEB 2003
NEWPORT

The Patent Office

1/77

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference — 200209039-1 GB

24FEB03 E787064-1 D01463
P01/7700 0.00-0304052.4

2. Patent application number *(The Patent Office will fill in this part)* — 0304052.4 — 22 FEB 2003

3. Full name, address and postcode of the or of each applicant *(underline all surnames)*

Hewlett-Packard Development Company, L.P.
20555 S.H. 249
Houston, TX 77070
USA

Patents ADP number *(if you know it)* — 855788600l

If the applicant is a corporate body, give the country/state of its incorporation — Texas ARB AIL 7/4/03

4. Title of the invention — Improved Streaming in Data Storage Drive

5. Name of your agent *(if you have one)*

"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*

David J. Marsh
Hewlett-Packard Ltd, IP Section
Filton Road, Stoke Gifford
Bristol BS34 8QZ

Patents ADP number *(if you know it)* — 828218800l

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and *(if you know it)* the or each application number

Country	Priority application number *(if you know it)*	Date of filing *(day / month / year)*

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application	Date of filing *(day / month / year)*

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if:*
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.
See note (d))

Yes

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	
Description	20
Claim(s)	7 CF
Abstract	1
Drawing(s)	6 + 6

10. If you are also filing any of the following, state how many against each item.

Priority documents	
Translations of priority documents	
Statement of inventorship and right to grant of a patent *(Patents Form 7/77)*	
Request for preliminary examination and search *(Patents Form 9/77)*	1
Request for substantive examination *(Patents Form 10/77)*	
Any other documents *(please specify)*	Fee Sheet

11. I/We request the grant of a patent on the basis of this application.

Signature David J. Marsh Date 20 Feb 03

David J. Marsh

12. Name and daytime telephone number of person to contact in the United Kingdom — Sue Holding Tel: 0117-312-9264

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.*

b) *Write your answers in capital letters using black ink or you may type them.*

c) *If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*

d) *If you have answered 'Yes' Patents Form 7/77 will need to be filed.*

e) *Once you have filled in the form you must remember to sign and date it.*

f) *For details of the fee and ways to pay please contact the Patent Office.*

IMPROVED STREAMING IN DATA STORAGE DRIVE

Field of the Invention

The present invention relates to the field of data storage, and particularly although not exclusively, to an improved method and apparatus for storage of data on a dynamic data storage medium.

Background to the Invention

It is known in the prior art to provide a data storage device externally to, or integrated within a host computer device, for the purpose of backing up data and systems stored on the host computer. Typically, data is sent from a host computer device to a tape drive unit, and the data is stored to a tape data storage medium provided in a cartridge which is removable from the tape drive unit.

Referring to Fig.1 herein, a known tape drive device 100 receives data from a host computer device 101, which may be networked to a plurality of other computers.

In general, data transferred from a host computer to a tape drive unit is 'bursty' that is, it is transmitted in chunks of data, followed by periods of no data. The data chunks are in general of variable and unpredictable length. Since the tape drive contains a tape transport mechanism which is electro-mechanical, and involves a tape data storage medium travelling past a read/write head, stopping and starting of the tape drive mechanism is best minimised or avoided for the following reasons.

- Firstly, excessive stopping and starting of the tape drive mechanism reduces the reliability of the mechanism over time.

- Secondly, stopping and starting of the tape drive mechanism requires re-positioning of the tape relative to the read/write head, which is time

consuming, and therefore reduces the rate at which data can be written to the tape data storage medium, particularly for linear tape drives.

In order to achieve optimum data throughput from the tape drive mechanism, the tape must be kept 'streaming', that is kept moving past the read/write head. To keep the tape data storage medium streaming past the read/write head, bursty data received from the host computer device is read into a buffer, which temporarily stores the data, removing some of the burstiness of the data. Continuous data exits the buffer at a more constant data rate determined by the rate at which the data can be written to the tape. In prior art tape drive devices, the rate determining step in writing data to a data storage medium is the relatively low rate at which data can be written from a write head to the tape. Although there is a problem of keeping the tape streaming when there are long bursts of no data arriving from the host, the existence of buffers helps to isolate the process of writing data to tape from the burstiness and drop outs in the incoming data stream from the host. However, if the average data rate received from the host computer drops below a rate at which the data continuously exists the buffer and is written to tape, then the buffer empties, and the tape must be stopped, repositioned to a position prior to a last data written, and then restarted once more data is available to fill the buffer of the tape drive unit.

To address the problem of tape stoppage and repositioning, there have been prior art systems developed which vary the speed of a tape past a write head, and thereby allow the tape to maintain streaming for a longer time period, without stoppage.

In US 6, 122,124, there is disclosed an adaptive tape speed method, in which the problem of tape stoppages is alleviated by keeping the tape moving past a write head at a reduced tape speed to match the incoming data rate, thus giving a slower data rate but with the advantage of maintaining streaming of the tape device.

However, in this prior art adaptive tape speed method, drive electronics limitations dictate a limited range of operation for the speed of the tape, which may not be sufficient to accommodate the full range in variations of data arriving from the host. Data arriving from the host may have variations in data rate which exceed the range of write data rates which correspond to the speeds available , and at which data can be written to tape.

Referring to Fig. 2 herein, there is illustrated schematically a prior art method of controlling tape speed by measuring buffer occupancy. A buffer device of a tape drive unit receives input data from a host computer, and produces an exit data stream which is output to a tape write head mechanism. The buffer has an occupancy level 201 of data stored in the buffer of between 0% and 100% of the full data storage capacity of the buffer. Depending upon the data rate of bursty data received from the host device in relation to the rate at which data exits the buffer, the occupancy level of the buffer can vary between 0% and 100%. Data arrives from the input host in bursts, fills up the buffer, and is output to the tape at a nominally constant data rate, which may be interrupted when the buffer becomes empty. Interruptions of the exit stream of data from the buffer cause tape stoppage and re-positioning.

In the prior art adaptive tape speed method, the occupancy level 201 of the buffer is electronically monitored, and used as a control signal for determining a speed of the tape past a write head.

Referring to Fig. 3 herein, there is illustrated schematically one example of a plot of tape speed past a write head, against time under various conditions of data received from a host computer device by a tape drive unit operating according to the known adaptive tape speed method.

Under normal operation, where data is being input from the host and filling up the buffer, and an instantaneous occupancy level of the buffer is above a first pre-determined limit, then the tape speed is controlled to be at its maximum value

300. However, if the data stream experienced from the host receives a drop out of data, then the buffer continues to empty of data, but no further data is received by the buffer and the instantaneous occupancy level falls. When the occupancy level falls to a second pre-determined limit, this triggers a reduction in tape speed to a second level 301. Since the tape drive has write electronics which matches the data rate of data exiting the buffer to the tape speed, there is a corresponding reduction in output data rate from the buffer. This keeps the tape streaming at a lower write data rate, until the buffer fills up again. If the buffer empties even further, then further pre-determined levels may trigger further reductions in tape speed 302. If the buffer becomes empty, then the tape must be stopped 303, which incurs the penalty of a time delay in repositioning the tape relative to the write head. Operation of the tape can resume at various tape speed levels, depending upon the amount of data received from the host and the occupancy level of the buffer.

Conventionally, during a data dump, or a data backup operation, host computers have generally been able to provide data to a tape drive unit at a higher average data rate than the data can be directly written to a tape data storage medium even though there may be instances of tape stoppage where bursty data is received from the host.

It is known to compress incoming data received from a host computer device prior to writing the data to a tape data storage medium. Data compression provides two main advantages as follows:

Firstly, it allows a greater amount of information to be stored to a tape storage medium using compression, than storing data directly in an uncompressed state to the tape data storage medium.

Secondly, since the rate of data arriving from the host computer is generally higher than the rate of data which can be written to tape, it allows a reduction in data rate written to the tape data storage medium, compared to the data rate

arriving from the host computer, the difference in data rate depending upon the amount of compression which can be applied to the data. This assists the tape drive in keeping up with writing the data to tape as it arrives from the host computer.

Referring to Fig. 4 herein, there is illustrated schematically components of another prior art host computer and tape drive unit, illustrating differences in data rates between data transferred from the host computer to the tape drive unit, and data written to tape. Data transferred from an internal data storage device 400 of a host computer 401 is transferred across a connection 402 in this example at a rate of 60 Mbytes/s. The data arrives at the tape drive unit 403, and is received by a data compression engine 404 which compresses the data. Varying compression ratios may be achieved depending upon the inherent compressibility of the incoming data. In the example shown, an average compression ratio of 2:1 is achieved, and data is written to the tape data storage medium at a data rate of 30 Mbytes/s.

In prior art host computer and tape drive units, data compression has a beneficial effect in at least partially isolating the data rate of data written to tape, from the bursty data, at a higher data rate, arriving from the host computer.

Some host computer operating systems provide a method of controlling the data compression ratio in the back up tape drive device via different device files and can disable compression if required. However, this is only done on a once and for all basis at the start of a data storage session and the compression ratio is not changed during the entire back up, regardless of ongoing performance of a data storage back up operation. The prior art host computer entity which can control compression has no visibility of the suitability of the data rate arriving at the tape drive entity which implements the compression, and so cannot set up the data compression ratio in the most effective way.

The adaptive tape speed method disclosed in US 6, 122, 124 does not disclose any data compression system, but rather deals with uncompressed data being written directly to tape.

However, there is a general trend to increase the write rate of data written from a write head to a tape in a tape drive unit, and as the write data rate from write head to tape increases to approach the data rate of incoming data from the host, the buffer system becomes less effective at isolating the write operation from drop outs in data arriving from the host, and the occurrence of stoppages increase, even where prior art methods such as the adaptive tape speed method are used. As the write data rate to tape increases towards the data transfer rate from host to the tape drive unit, the problem of stoppages becomes more acute with a higher incidence of tape stoppages occurring.

As tape drives get faster, and are capable of writing data to tape at a higher data rate, they do not necessarily represent the rate determining stage in a data storage system when performing data storage operations, for example data back ups. System performance may typically be limited by the ability of a host computer to supply data fast enough to keep the tape drive streaming. If the incoming data rate from a host computer drops below a minimum acceptable data rate, then the tape must be stopped, repositioned prior to the last data written, and then restarted once sufficient data is available from the host computer.

Once the host stops supplying data for an extended period, and streaming of the tape stops, a delay of several seconds is incurred whilst the tape repositions itself, which is a far higher delay than a latent delay caused by the host in recommencing supply of bursts of data. Therefore, stoppages in streaming are to be avoided where ever possible, since the stoppages become the rate determining step in transfer of data from the host of the tape data storage medium when they occur.

Summary of the Invention

According to a first aspect of the present invention there is A tape drive unit comprising: a data compression engine capable of applying compression to an incoming data stream; a buffer memory capable of storing data of said incoming data stream; a monitoring module capable of monitoring a data occupancy level of said buffer memory; and a control module capable of enabling or disabling said data compression engine; wherein said control module is operable to enable or disable said data compression engine depending upon said data occupancy level of said buffer memory.

According to a second aspect of the present invention, there is provided a data processing device comprising: a data compression engine capable of applying compression to an incoming data stream; a buffer memory capable of storing data of said incoming data stream; wherein said device being operable to enable or disable said data compression engine depending upon said data occupancy level of said buffer memory.

According to a third aspect of the present invention, there is provided a data processing device comprising: means for applying compression to an incoming data stream; and means for storing data of said incoming data stream; wherein said data processing device is operable to enable or disable the data compression means depending upon a data occupancy level of the data storage means.

According to a fourth aspect of the present invention there is provided a data processing device comprising: a data compression engine capable of applying compression to a data stream received by said device; and a buffer memory capable of storing data of said received data stream; said device being operable such that said compression engine is controlled to apply compression processing to said received data stream in response to said amount of data in said buffer memory being at a relatively high value, and is controlled to be inactive such that said received data stream is entered into said buffer memory

without being compression processed by said compression engine, in response to said amount of data in said buffer memory being at a relatively low value.

According to a fifth aspect of the present invention, there is provided program instructions for causing a data processor to: monitor a data occupancy level of a buffer memory; and enable or disable a compression engine, depending upon a data occupancy level of said buffer memory.

According to a sixth aspect of the present invention, there is provided program instructions for causing a data processor to: monitor a data occupancy level of a buffer memory; enable or disable a compression engine, depending upon a data occupancy level of said buffer memory; said program instructions being stored in at least one of a volatile or a non volatile memory.

According to a seventh aspect of the present invention, there is provided program instructions for controlling a processor to apply data processing to an incoming data stream, said program code comprising: a monitoring module capable of monitoring a data occupancy level of a buffer memory; and a control module capable of enabling or disabling a compression engine; wherein said control module is operable to send an enable signal or a disable signal for enabling or disabling said data compression engine, depending upon said data occupancy level of said buffer memory monitored by said monitoring module.

According to an eighth aspect of the present invention, there is provided a method of writing data to a tape data storage medium, said method comprising: receiving an incoming data stream into a buffer memory; in response to said data occupancy level being at a pre-determined threshold level, applying data compression to said incoming data stream, upstream of said buffer.

According to a ninth aspect of the present invention, there is provided a method of writing data to a tape data storage medium, said method comprising: receiving an incoming data stream into a buffer memory; monitoring a data

occupancy level of said buffer memory; in response to said data occupancy level being at a pre-determined threshold level, applying data compression to said incoming data stream, upstream of said buffer; and in response to said data occupancy level falling below said pre-determined threshold level, receiving said incoming host data stream directly into said buffer, without applying any compression.

According to a tenth aspect of the present invention there is provided a method of writing data to a tape data storage medium, said method comprising: receiving a data stream in bursts of data at a first data rate into a buffer memory; transferring data of said data stream from said buffer memory to said tape data storage medium at a second data rate; monitoring a data occupancy level of said buffer memory; depending upon a value of said data occupancy level, varying a speed of said tape data storage medium past a write head; and depending upon a value of said data occupancy level, activating or de-activating a data compression engine positioned up stream of said buffer memory, for compressing data of said host data stream prior to entry into said buffer memory.

According to an eleventh aspect of the present invention, there is provided a method of controlling a data rate of data exiting a buffer memory, said method comprising: receiving a data stream at a first data rate: selectively applying or not applying compression to said data stream to produce a compression processed data stream; receiving said selectively compression processed data stream into a buffer memory; outputting said selectively compression processed data stream as an output of said buffer memory at a second data rate.

According to a twelfth aspect of the present invention, there is provided a method of data processing a stream of data comprising: receiving said data stream; entering data of said received data stream in a buffer memory; applying compression processing to said received data stream under conditions of said amount of data in said buffer memory being relatively high; and entering said received data stream into said buffer memory without compression processing

said data stream under conditions of said amount of data in said buffer memory being relatively low.

Other aspects are as recited in the claims herein.

Brief Description of the Drawings

For a better understanding of the invention and to show how the same may be carried into effect, there will now be described by way of example only, specific embodiments, methods and processes according to the present invention with reference to the accompanying drawings in which:

Fig. 1 illustrates schematically a prior art host computer and prior art tape drive unit capable of receiving data from the host computer;

Fig. 2 illustrates schematically an occupancy level of a buffer comprising the prior art tape drive unit;

Fig. 3 illustrates schematically a plot of tape speed against time for a prior art tape drive unit using an adaptive tape speed method;

Fig. 4 illustrates schematically components of a further prior art host computer and prior art tape drive unit, illustrating a difference between a data rate between the host computer and the tape drive, and a write data rate to a tape data storage medium;

Fig. 5 illustrates schematically modules of a tape drive unit according to a specific implementation of the present invention;

Fig. 6 illustrates schematically a set of functional modules for monitoring buffer occupancy level, controlling a compression engine, and controlling a tape transport mechanism to apply an adaptive tape speed control method according to a specific implementation of the present invention;

Fig. 7 illustrates schematically a first method operation of a tape drive unit, according to a specific method of the present invention, for regulating a stream of data between a host interface and a tape data storage medium; and

Fig. 8 illustrates schematically a second method of operation of a tape drive unit according to the specific method of the present invention for regulating a flow of data between a host interface and a tape data storage medium, where the method operates in conjunction with a prior art adaptive tape speed method.

Detailed Description of a Specific Mode for Carrying Out the Invention

There will now be described by way of example a specific mode contemplated by the inventors. In the following description numerous specific details are set forth in order to provide a thorough understanding of the embodiments of present invention. It will be apparent however, to one skilled in the art, that the present invention may be practiced without limitation to these specific details. In other instances, well known methods and structures have not been described in detail so as not to unnecessarily obscure the description.

Referring to Fig. 5 herein, there is illustrated schematically a tape drive unit according to a specific embodiment of the present invention. The tape drive unit 500 comprises a host interface 501 for receiving incoming data from a host computer; a data compression engine 502 for compressing data received from the host interface; a buffer 503 for transiently storing data; a read/write control component 504 for controlling read and write operations to and from a set of read/write heads 505; a processor 506 and an associated set of control elements 507 for providing instructions to the processor; a set of drive motors 508 for driving a tape cassette transport mechanism 509 which drives a tape cassette 510; and a servo system 511 for controlling the set of drive motors.

The control elements 507 may be implemented as "firmware" modules comprising a set of program instructions, stored in a non volatile memory such as

an EEPROM. Alternatively the processor 506 could be dispensed with and the elements 507 implemented by a logic circuit for example an application specific integrated circuit or a programmable logic array. Alternatively, any other convenient arrangement could be used to implement the control elements 507

A data path through the tape drive mechanism exists between the host interface, data compression engine, buffer, read/write electronics and read/write heads.

The processor 506 and control elements 507 control the host interface, data compression engine, buffer, read write electronics 504, and servo system 511 to write data to the tape data storage medium via the write heads, and conversely, to read data from the tape data storage medium via the read heads.

According to a specific implementation, the control elements operate to monitor both the incoming data rate from the host computer device, and also the compression ratio being achieved by the data compression engine 502 within the tape drive. These two pieces of information can be used to determine when to turn on or off data compression within the tape drive unit. When data compression is turned off, this effectively reduces the write data rate required to maintain continuous streaming of the tape data storage medium, compared to when compression is applied.

According to a specific implementation, control of the data compression ratio is achieved at the point where the write data rate to the tape is known, that is, within the tape drive unit itself.

Referring to Fig. 6 herein, control elements 507 comprise the following functional modules; a buffer occupancy level monitoring module 601; a data compression engine control module 603; and optionally, an adaptive tape speed control module 603, in addition to other known modules for controlling the read write electronics, and servo-system. The other known control modules may

comprise a known adaptive tape speed control module 603, for embodiments where the adaptive tape speed method is used in conjunction with the novel method disclosed herein.

There will now be described a first method of operation of a tape drive unit in which a data compression operation is controlled depending upon a data occupancy level of a buffer in a tape drive unit.

Referring to Fig. 7 herein, there is illustrated schematically processes comprising a first specific method according to the present invention for controlling a streaming operation of a tape data storage medium, by varying compression of an incoming host data stream.

In the following description, where compression is referred to, this means compression by the compression engine at the tape drive unit. The skilled person will appreciate that the host data received from the host computer may have been pre-compressed to a varying extent by an application at the host computer, or may have been compressed elsewhere in a network. The tape drive unit is capable of applying further compression processing, with a varying compression ratio, depending upon the inherent compressibility of the data being received from the host computer.

In process 700, the buffer is continuously monitored to check its occupancy level. A pre-set occupancy level is set as a trigger level to trigger enabling of the data compression engine. As long as the occupancy level of the buffer is maintained above that trigger level, then the data compression engine continues to compress incoming data received from the host device, so that the data received by the buffer, down stream of the data compression engine, is compressed data. Various levels of compression may be applied by the compression engine, depending upon the intrinsic compressibility of the incoming data from the host device. Where the host device sends to the tape drive unit data which has already undergone a high level of compression, it may be that the

data cannot be further compressed. On the other hand where there is a large amount of redundancy of information within the incoming host data, then a higher degree of compression may be achieved. The compression engine continues to be enabled, as long as the occupancy level of the buffer remains above the pre-set trigger level. However, if the occupancy level of the buffer falls below the pre-set trigger level, then in process 701, the data compression is disabled by the compression engine controller module 602 at the next available point at which the compression can be turned off. Because the compression engine takes chunks of data in order to run its compression algorithm, it may be that the compression engine may be already processing a chunk of data, and so the compression engine cannot be disabled immediately, until it has finished compressing that whole chunk of data. As soon as the compression engine finishes compressing its current data chunk, consistent with proper operation of the compression algorithm and the tape format specification, then the compression engine is disabled, and the host data is fed into the buffer, filling up the buffer without any compression having been applied at the tape drive unit. With compression processed or non-compression processed data entering the buffer, data is continuously streamed out of the buffer through the read/write electronics and to the write heads, to be written to the tape data storage medium. During transfer of data which has not been compression processed by the compression engine to the buffer, as an ongoing process, the buffer occupancy level is continuously monitored in process 702. If the buffer occupancy level fills up above the trigger level again, then in process 703 the data compression engine is re-enabled at the next available point in the data, and processes 700 – 702 continue. In this state, data is arriving from the host computer at a high enough rate that the data can be compressed at the tape drive, and fills up the buffer, faster than the rate of emptying of the buffer and faster than the rate of writing the data being emptied from the buffer directly on to the tape.

Whilst the data compression engine is disabled, as long as the buffer occupancy level remains below the trigger level for re-enabling the data compression engine in process 702, then in process 704 it is continuously

monitored whether the incoming data rate from the host computer is above a minimum acceptable data rate. If the data rate is above a minimum acceptable data rate, then in process 705 an amount of non-usage of the data compression engine which has occurred is monitored and it is checked whether that amount of non-usage has extended beyond the pre-determined limit. Non-usage of the compression engine can be determined either as a time period over which the compression engine has been disabled, or as a number of bytes which have been received from the incoming host computer whilst the compression engine has been disabled, or an amount of data which has been written to tape whilst the compression engine has been disabled. As long as the pre-determined limit for non-usage of the compression engine has not been reached, then the compression engine remains disabled, and incoming data from the host which has not been compression processed by the compression engine continues to be fed directly into the buffer. At the same time, as an ongoing process, data is fed out of the buffer to be written directly to the tape, which continues to keep streaming. The servo system is controlled to operate the motors, such that the transport mechanism to the tape data storage medium keeps the tape moving past the write heads as data is written to the tape, even when data which has not been compression processed by the compression engine is being written.

However, once the non-usage limit of the compression engine has been reached, either the pre-set time limit is reached, or a pre-determined number of bytes have been received from the host, or written to tape, and the buffer occupancy level is still below the trigger level, then the tape transport is stopped and the tape is re-positioned. This stoppage of tape streaming can incur a time delay in re-positioning the tape of a few seconds. In process 706, the tape streaming is stopped and the tape is re-positioned. At the same time, the data compression engine maybe re-enabled, and once the tape begins to stream again, operation commences from process 700 as described herein above.

On the other hand, in process 704, where the buffer occupancy level is below the trigger level, if at any time the average host data rate drops below the

pre-set minimum acceptable data rate, indicating that not enough data is being received from the host computer, then stoppage of the tape occurs immediately and tape re-positioning occurs, which takes a few seconds. At the same time, the compression engine may be re-enabled in process 706, and further operation continues from process 700 substantially as described herein above.

In the above operation, by monitoring the data occupancy level of the buffer relative to a pre-determined trigger level, when the occupancy level falls below that pre-determined trigger level, the compression engine is disabled, so that data received from the host is fed directly into the buffer at a data rate of approximately the same as the rate at which the data is received from the host computer, and can therefore start filling up the buffer since in general this data rate will be faster than the data rate at which data leaves the buffer to be written to the tape. Once the data in the buffer fills to above the pre-determined trigger level, then the compression engine is re-activated so that incoming data from the host computer is compression processed at the tape drive, resulting in a relatively lower rate of data being input into the buffer, than the incoming host data rate. In this mode of operation, data is being compressed once it arrives at the tape drive unit, and the rate determining stage in the tape drive unit is the rate at which data can be written directly to the tape.

There will now be described a second specific method of operation of the tape drive unit of Figs. 5 and 6 herein, in which the tape drive unit operates the known prior art adaptive tape speed method as described herein above, in combination with selective application of data compression to an incoming data stream, according to the specific implementation of the present invention.

Referring to Fig. 8 herein, there is illustrated schematically process steps according to a second specific method of the present invention.

In process 800, it is checked whether the buffer occupancy level is above a lower threshold limit for the adaptive tape speed method to operate. Whilst the

buffer occupancy level is above this level, there may be other higher occupancy threshold levels used by the adaptive tape speed method. Depending upon the buffer occupancy level, above the lower threshold limit the tape is speeded up or slowed down, in order to avoid stoppages in the movement of the tape. The lower threshold limit of the adaptive tape speed method represents the lowest buffer occupancy at which the adaptive tape speed method will operate, and if the buffer occupancy drops below this level, then the tape speed cannot be slowed any further. In process 801, if the buffer occupancy level falls below the lower threshold level for the adaptive tape speed method, then the data compression engine is disabled at the next available point. As explained herein above, the compression engine may be unable to be immediately disabled, because it may be in the middle of processing a data chunk. Therefore, the present data chunk being compression processed must be completed, before the data compression can be disabled. In process 802, once the compression engine has been disabled, the buffer occupancy level is continuously monitored relative to the lower ATS threshold level. If the buffer occupancy level rises above the lower ATS threshold level again, then in process 803 the data compression is re-enabled at the next opportunity, and normal operation of the tape drive unit, within the ATS mode can continue as described herein above with reference to processes 800-802. However, if the buffer occupancy level remains below the ATS threshold limit, then in process 804 it is continuously checked whether the incoming host data rate is above a minimum acceptable rate. As long as the data rate is above a minimum acceptable data rate, then in process 806 it is continuously monitored how much non-usage of the compression engine has occurred. This can be monitored either as a pre-set limit during which the compression engine is disabled, or a number of bytes which have been received from the host computer whilst the compression engine has been disabled, or a number of bytes written to tape whilst the compression engine is disabled. Provided the non-usage of the data compression engine is within a pre-determined limit, then the tape drive continues in its mode of operation where the data compression engine is turned off, and the adaptive tape speed method is operating on its lowest tape speed. In this mode of operation data continues to

be written to the tape, and the tape keeps streaming at its lowest speed. The data entering the buffer is not compression processed within the tape drive during this mode.

If, whilst the buffer occupancy level is below the lower ATS threshold, the incoming host data rate falls below a minimum acceptable data rate, then in process 805 the tape transport mechanism is stopped immediately and the tape is re-positioned. Whilst this occurs, further data may arrive from the host computer, the compression engine is re-enabled, and further operation continues as in processes 800-801 as described herein above.

In the second specific method, a combination of the prior art adaptive tape speed method, and control of the compression engine is used to give an extended operating range during which tape transport can continue without stoppages of streaming. The adaptive tape speed method allows for variation of tape speed and slowing down of the tape speed to continue to write compressed data as long as there is a sufficient amount of data arriving from the host computer. However, where the rate of data arriving from the host drops to such a level that the buffer starts to empty below an occupancy level at which the adaptive tape speed method would normally cause a tape stoppage, the compression engine is disabled, which effectively increases the rate of data which is fed from the host interface directly to the buffer, compared to the situation where compression is being applied. This enables the buffer to fill up relatively faster than if compression of data where being applied by the compression engine. Continuously with this process, data exits the buffer at a data rate consistent with the lowest tape speed. As soon as the rate of incoming data from the host computer is high enough that the buffer occupancy level raises above the lower ATS threshold, then the compression engine is re-enabled and compression of incoming data resumes.

The specific method disclosed herein may be implemented in a tape drive unit in addition to the known adaptive tape speed method, or may be implemented independently of the adaptive tape speed method.

As a data rate approaches the lower limit of the adaptive tape speed range, provided that the current compression ratio being applied is above a certain threshold level (which can be determined empirically), then the data compression engine is disabled by the controller.

This state remains until either:

a) The incoming data rate recovers to above the lower limit of the adaptive tape speed threshold level; or

b) The incoming data rate drops below the minimum data rate of the drive; or

c) The compression engine controller determines that the drive has been in this state for too long.

In case (a) above, the controller would then re-enable compression, and operation of the tape drive would continue.

In case (b) there is nothing else which can be done to prevent stoppage of the tape. The drive will have to stop and re-position the tape. Whether to re-start the tape with compression enabled or disabled, is an implementation specific choice.

In case (c) this is a means for preventing the drive running in a mode in which data is not being compressed, for the whole duration of a back up. If data is continuously written to the tape in un-compressed mode, this leads to reduced tape data storage capacity.

One possibility for 'tuning' the system would be allow the host to specify, at the start of the back up session, the maximum time or amount of data to be written in an uncompressed mode. Also, the host computer could select whether to set the tape drive up to 'optimise for performance' i.e. optimise for keeping the tape transport mechanism streaming, or to 'optimise for capacity' in which case, the above specific methods of operation of a tape drive unit would not be enabled, and the tape drive would operate in a mode in which compression is always applied, or in a mode with the adaptive tape speed feature, in which compression is always applied, but at the disadvantage of continuous interruptions through tape stoppage and re-positioning, with all the associated problems of tape drive reliability of sub-optimal delays due to tape stoppage and re-positioning.

Specific implementations as described herein address the issue of stoppages of tape streaming in a tape drive unit.

Claims:

1. A tape drive unit comprising:

a data compression engine capable of applying compression to an incoming data stream;

a buffer memory capable of storing data of said incoming data stream;

a monitoring element capable of monitoring a data occupancy level of said buffer memory; and

a control element capable of enabling or disabling said data compression engine;

wherein said control element is operable to enable or disable said data compression engine depending upon said data occupancy level of said buffer memory.

2. The tape drive unit as claimed in claim 1, comprising:

a tape transport mechanism for transporting a tape data storage medium past a set of write heads;

wherein said tape transport mechanism is operable to continue streaming of said tape, whilst said data compression engine is in an enabled mode, and whilst said compression engine is in a disabled mode.

3. The tape drive unit as claimed in claim 1, comprising:

a tape transport mechanism for transporting a tape data storage medium past a set of write heads; and

a tape speed control element for controlling said tape transport mechanism for transporting said tape at a variable speed:

a said speed of tape is selectable according to a data occupancy level of said buffer memory.

4. A data processing device comprising:

a data compression engine capable of applying compression to an incoming data stream;

a buffer memory capable of storing data of said incoming data stream;

said device being operable to enable or disable said data compression engine depending upon said data occupancy level of said buffer memory.

5. The data processing device as claimed in claim 4, further comprising:

a monitoring element capable of monitoring a data occupancy level of said buffer memory; and

a control element capable of enabling or disabling said data compression engine.

6. A data processing device comprising:

means for applying compression to an incoming data stream; and

means for storing data of said incoming data stream;

wherein said data processing device is operable to enable or disable the data compression means depending upon a data occupancy level of the data storage means.

7. A data processing device comprising:

a data compression engine capable of applying compression to a data stream received by said device; and

a buffer memory capable of storing data of said received data stream;

said device being operable such that said compression engine is controlled to apply compression processing to said received data stream in response to said amount of data in said buffer memory being at a relatively high value, and is controlled to be inactive such that said received data stream is entered into said buffer memory without being compression processed by said compression engine, in response to said amount of data in said buffer memory being at a relatively low value.

8. Program instructions for causing a data processor to:

monitor a data occupancy level of a buffer memory;

enable or disable a compression engine, depending upon a data occupancy level of said buffer memory.

9. Program instructions for causing a data processor to:

monitor a data occupancy level of a buffer memory;

enable or disable a compression engine, depending upon a data occupancy level of said buffer memory;

said program instructions being stored in at least one of a volatile or a non volatile memory.

10. Program instructions for controlling a processor to apply data processing to an incoming data stream, said program code comprising:

a monitoring module capable of monitoring a data occupancy level of a buffer memory; and

a control module capable of enabling or disabling a compression engine;

wherein said control module is operable to send an enable signal or a disable signal for enabling or disabling said data compression engine, depending upon said data occupancy level of said buffer memory monitored by said monitoring module.

11. A method of writing data to a tape data storage medium, said method comprising:

receiving an incoming data stream into a buffer memory;

monitoring a data occupancy level of said buffer memory;

in response to said data occupancy level being at a pre-determined threshold level, applying data compression to said incoming data stream, upstream of said buffer.

12. The method as claimed in claim 11, further comprising:

in response to said data occupancy level falling below said pre-determined threshold level, inputting said incoming data stream directly into said buffer, without applying any compression.

13. A method of writing data to a tape data storage medium, said method comprising:

receiving an incoming data stream into a buffer memory;

monitoring a data occupancy level of said buffer memory;

in response to said data occupancy level being at a pre-determined threshold level, applying data compression to said incoming data stream, upstream of said buffer; and

in response to said data occupancy level falling below said pre-determined threshold level, receiving said incoming host data stream directly into said buffer, without applying any compression.

14. A method of writing data to a tape data storage medium, said method comprising:

receiving a data stream in bursts of data at a first data rate into a buffer memory;

transferring data of said data stream from said buffer memory to said tape data storage medium at a second data rate;

monitoring a data occupancy level of said buffer memory;

depending upon a value of said data occupancy level, varying a speed of said tape data storage medium past a write head; and

depending upon a value of said data occupancy level, activating or de-activating a data compression engine positioned up stream of said buffer memory, for compressing data of said host data stream prior to entry into said buffer memory.

15. A method of controlling a data rate of data exiting a buffer memory, said method comprising:

receiving a data stream at a first data rate:

selectively applying or not applying compression to said data stream to produce a compression processed data stream;

receiving said selectively compression processed data stream into a buffer memory;

outputting said selectively compression processed data stream as an output of said buffer memory at a second data rate.

16. The method as claimed in claim 15, wherein said process of selectively applying compression to said data stream comprises:

applying a compression algorithm to said data stream in response to a data occupancy level of said buffer memory being at a pre-determined threshold limit; and

disabling application of said compression algorithm to said data stream, in response to a data occupancy level of said buffer memory falling below said pre-determined threshold limit.

17. The method as claimed in claim 15 or 16, wherein said output of said buffer memory at said second data rate is written to a tape data storage medium.

18. A method of data processing a stream of data comprising:

receiving said data stream;

entering data of said received data stream in a buffer memory;

applying compression processing to said received data stream under conditions of said amount of data in said buffer memory being relatively high; and

entering said received data stream into said buffer memory without compression processing said data stream under conditions of said amount of data in said buffer memory being relatively low.

Abstract

IMPROVED STREAMING IN DATA STORAGE DRIVE

A tape drive unit comprising: a data compression engine capable of applying compression to an incoming data stream; a buffer memory capable of storing data of said incoming data stream; a monitoring module of monitoring a data occupancy level of said buffer memory; and a control module capable of enabling or disabling said data compression engine; wherein said control module is operable to enable or disable said data compression engine depending upon said data occupancy level of said buffer memory.

Fig. 7

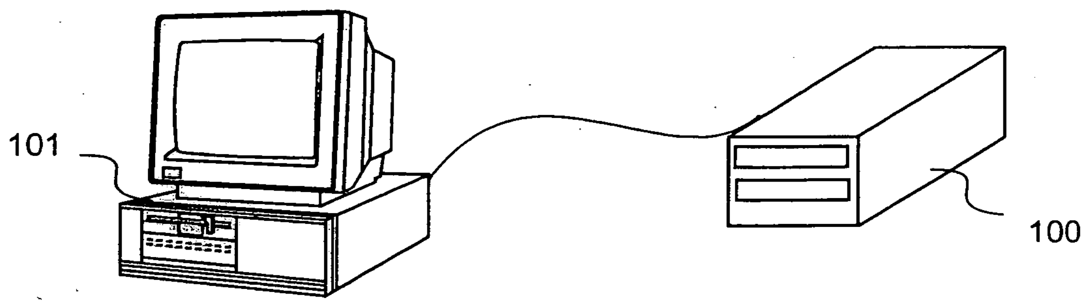

Fig. 1
(Prior Art)

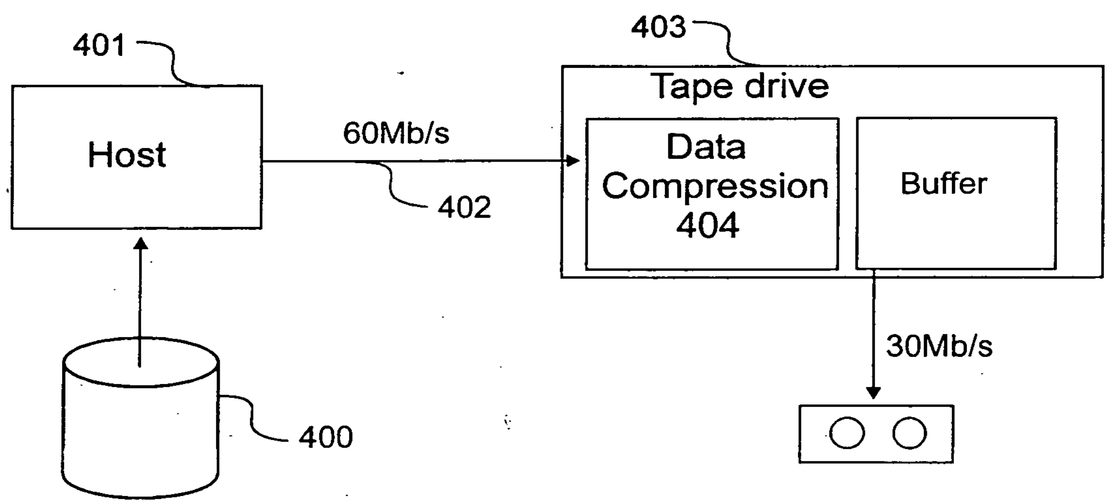

Fig. 4
(Prior Art)

THIS PAGE BLANK (USPTO)

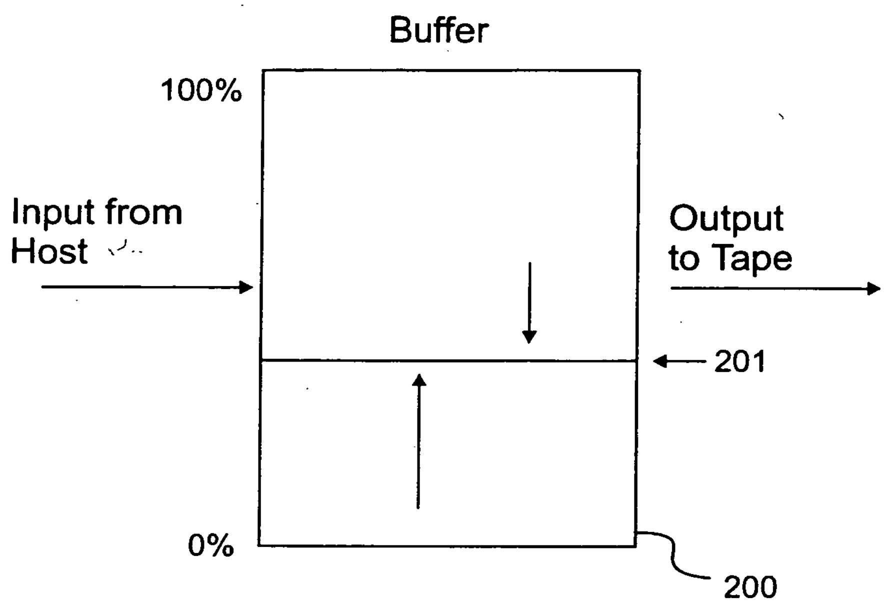

Fig. 2
(Prior Art)

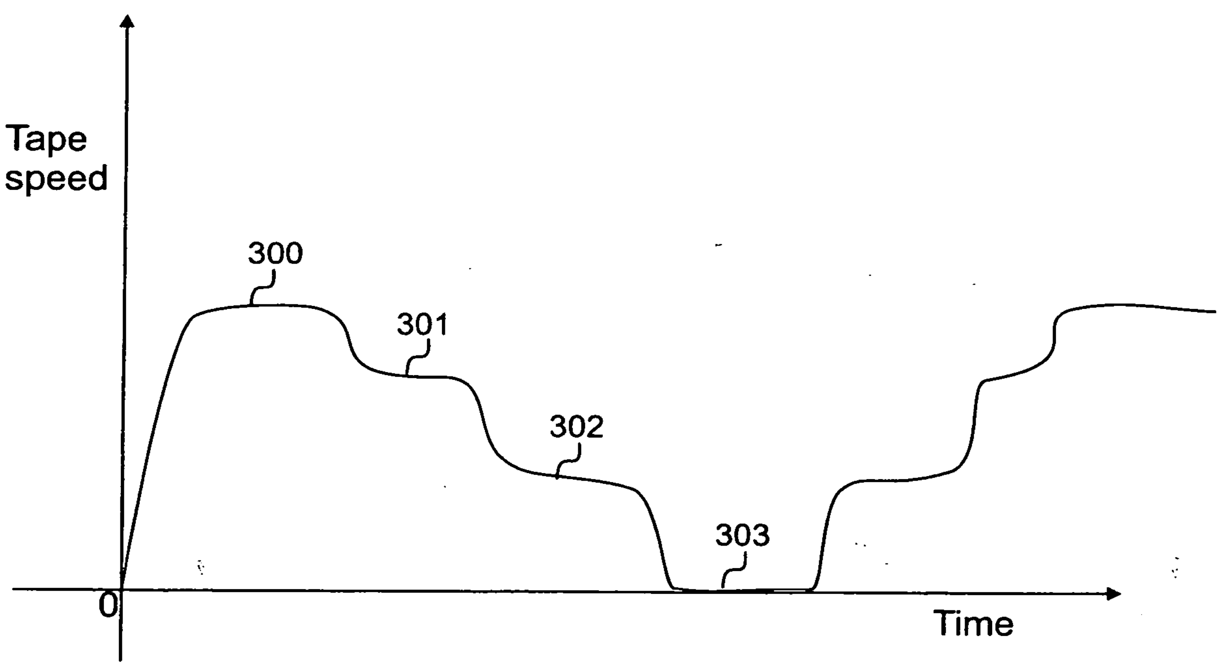

Fig. 3
(Prior Art)

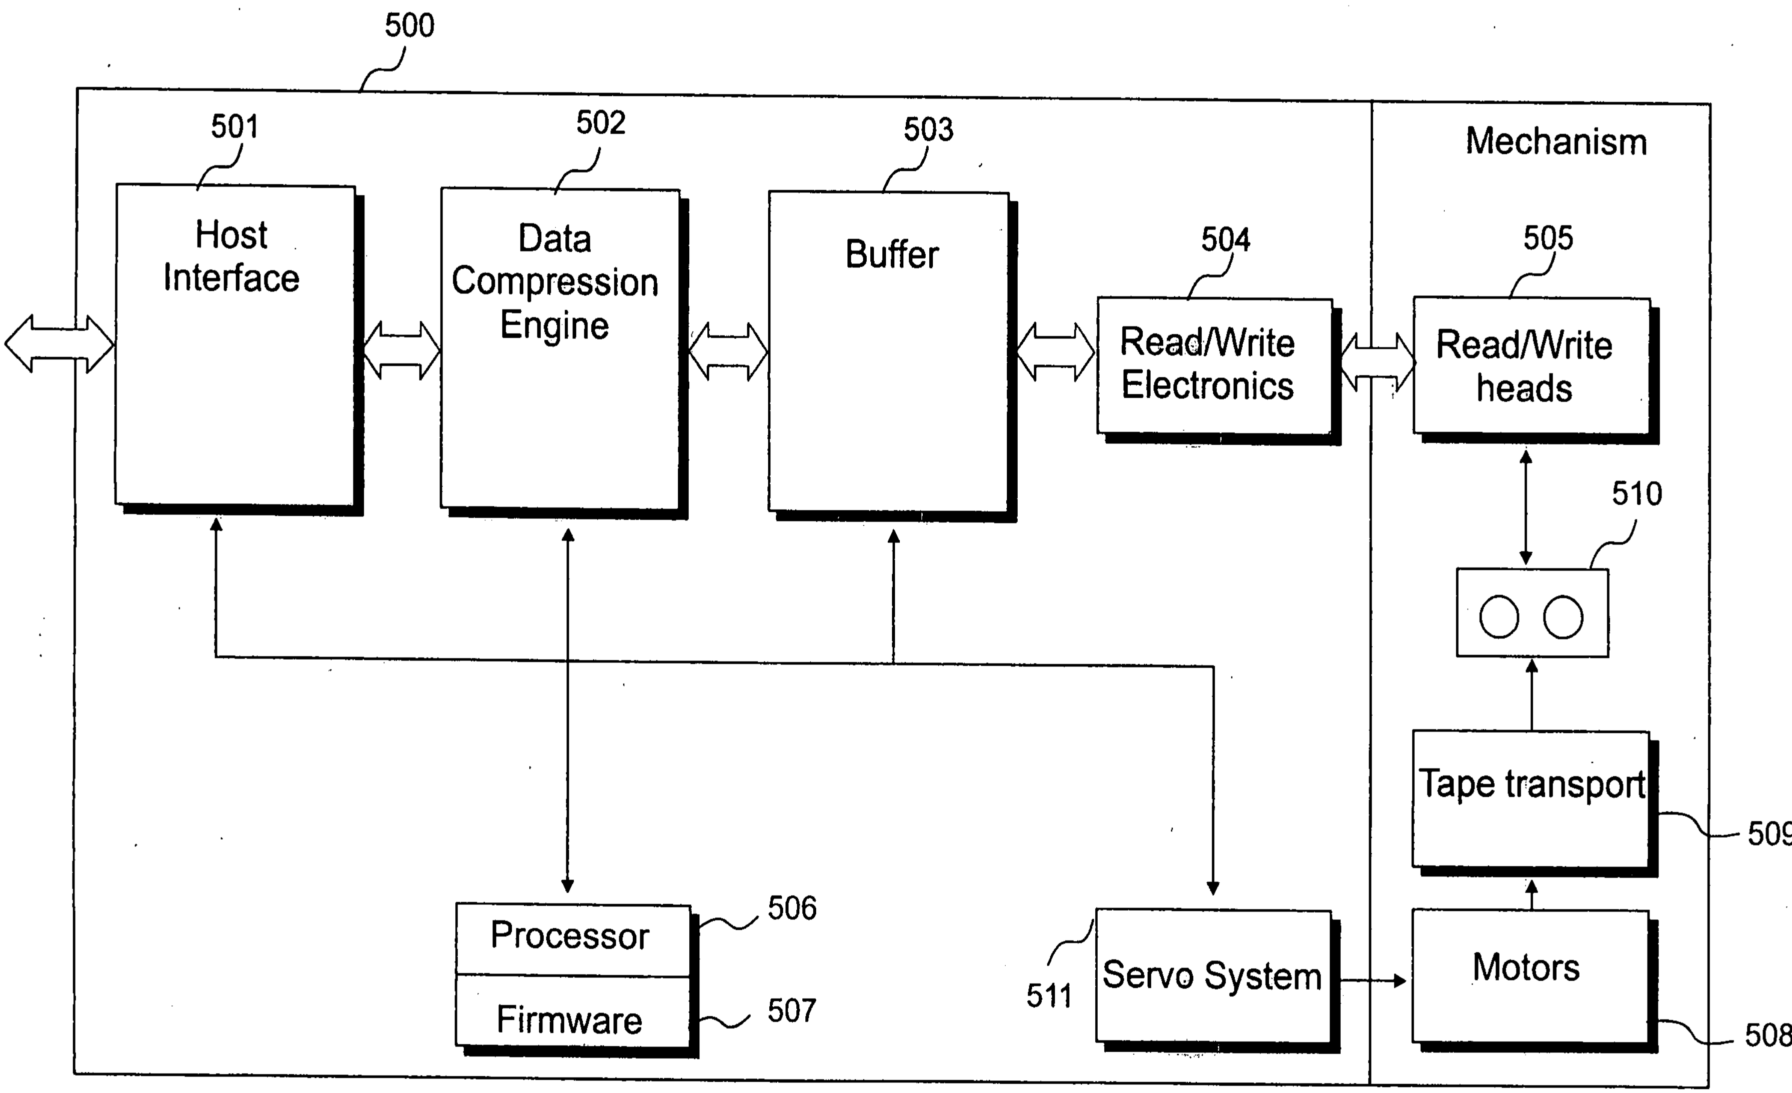
500
501
Host Interface
502
Data Compression Engine
503
Buffer
504
Read/Write Electronics
Mechanism
505
Read/Write heads
510
Tape transport
509
Motors
508
Processor
506
Firmware
507
511
Servo System

Fig. 5

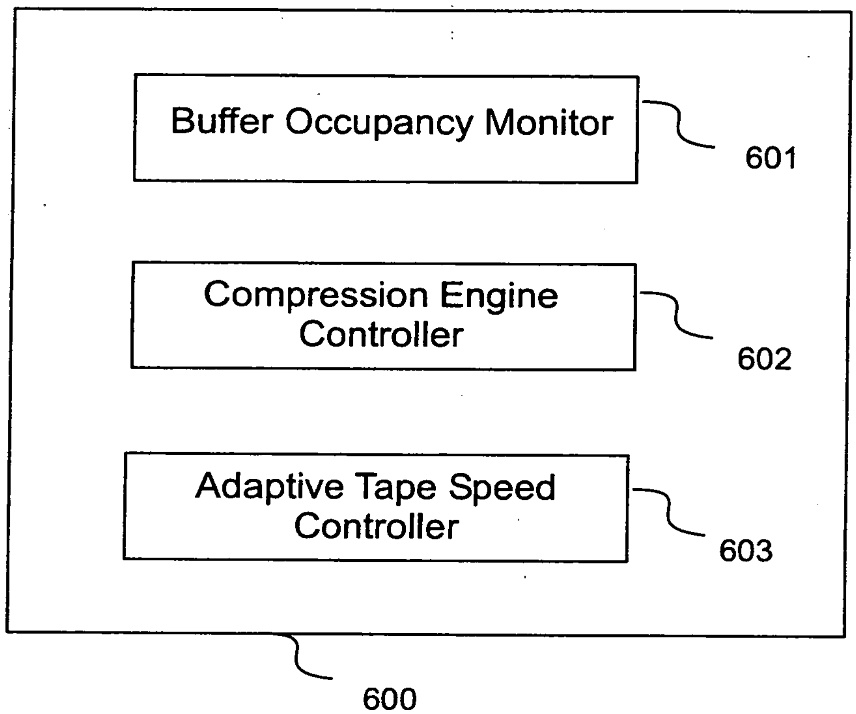
Buffer Occupancy Monitor
601
Compression Engine Controller
602
Adaptive Tape Speed Controller
603
600

Fig. 6

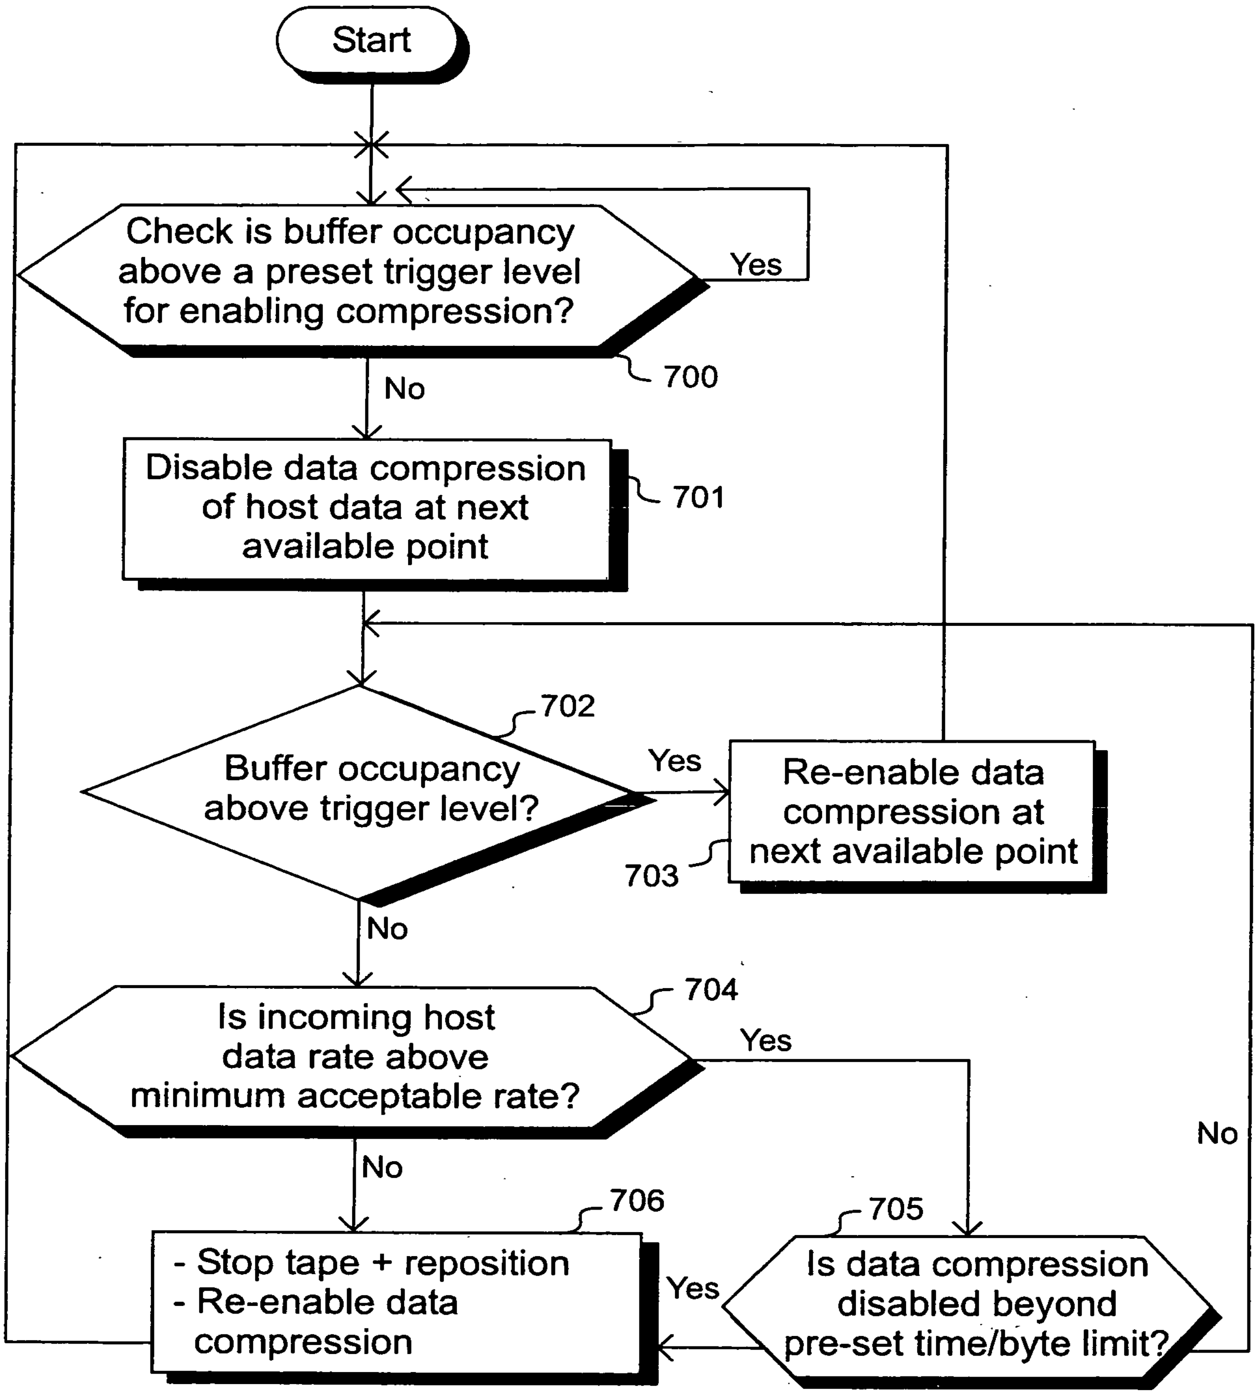

Start
Check is buffer occupancy above a preset trigger level for enabling compression?
Yes
No
700
Disable data compression of host data at next available point
701
702
Buffer occupancy above trigger level?
Yes
Re-enable data compression at next available point
703
No
704
Is incoming host data rate above minimum acceptable rate?
Yes
No
No
706
705
- Stop tape + reposition
- Re-enable data compression
Yes
Is data compression disabled beyond pre-set time/byte limit?

Fig. 7

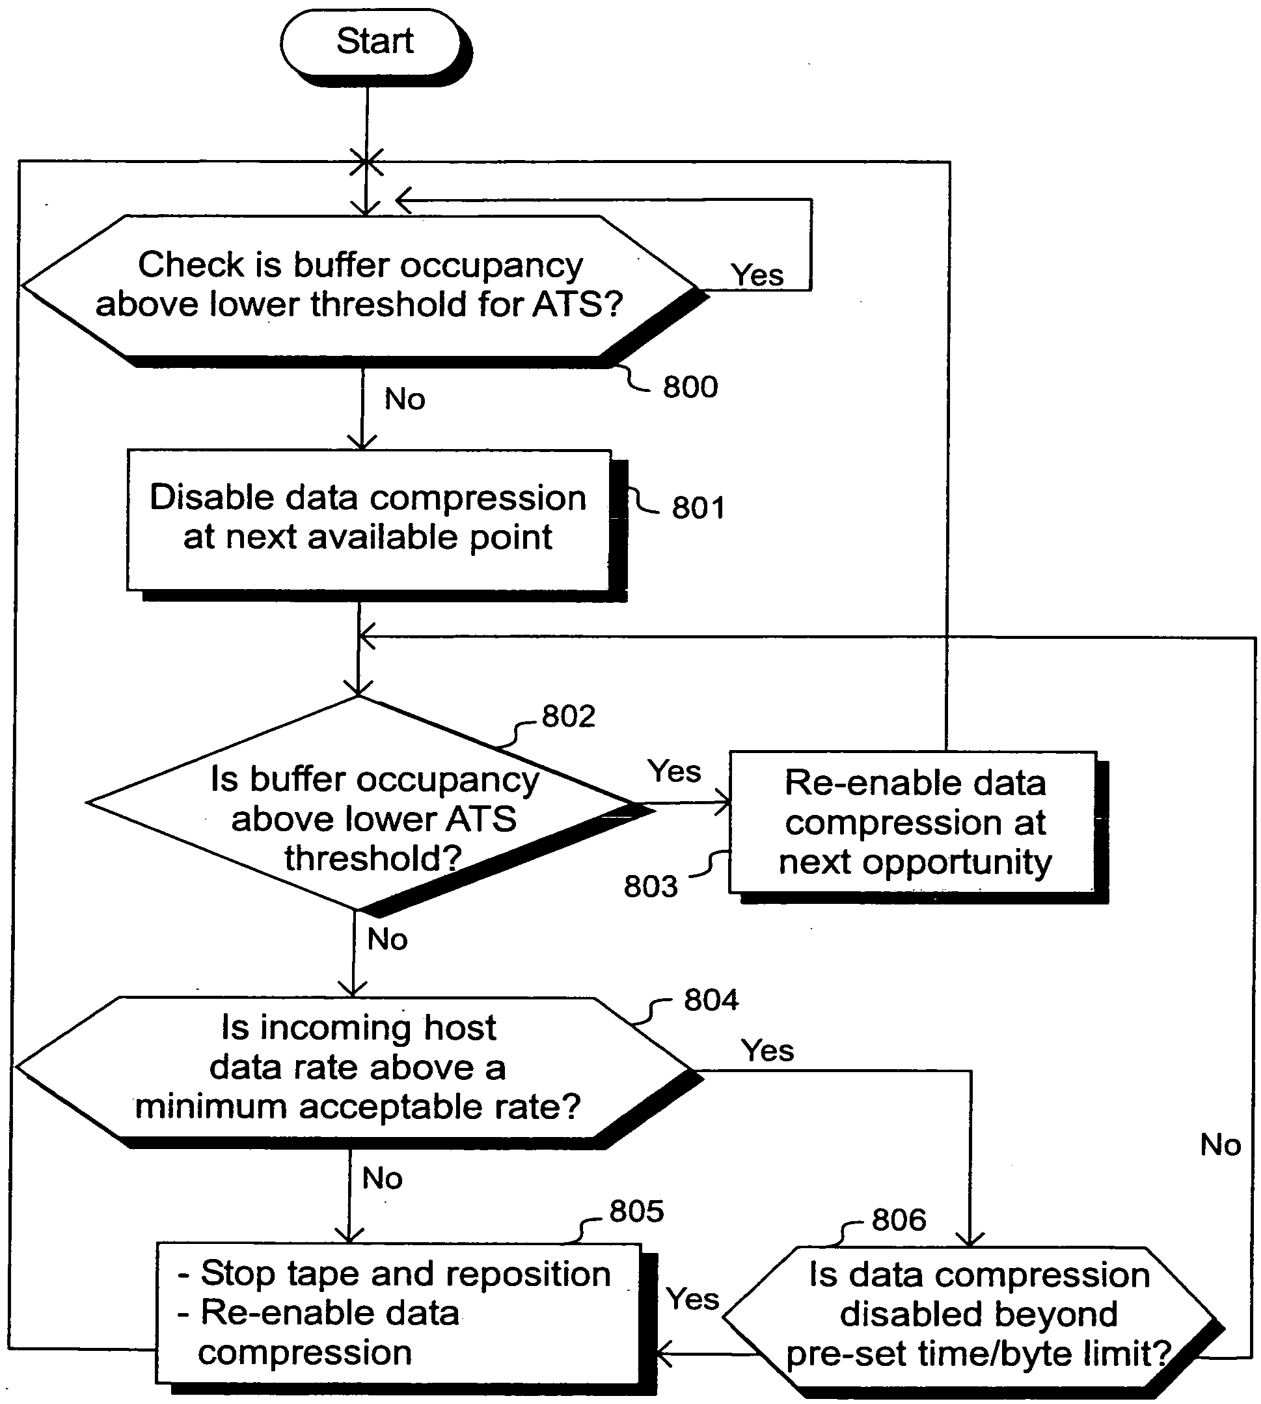
Start
Check is buffer occupancy above lower threshold for ATS?
Yes
No
800
Disable data compression at next available point
801
802
Is buffer occupancy above lower ATS threshold?
Yes
Re-enable data compression at next opportunity
803
No
804
Is incoming host data rate above a minimum acceptable rate?
Yes
No
No
805
806
- Stop tape and reposition
- Re-enable data compression
Yes
Is data compression disabled beyond pre-set time/byte limit?

Fig. 8